# *SaLF*: Sparse Local Fields for Multi-Sensor Rendering in Real-Time

Yun Chen*, Matthew Haines*, Jingkang Wang, Sahil Jain, Krzysztof Baron-Lis,
Sivabalan Manivasagam, Ze Yang, Raquel Urtasun

waabi

https://waabi.ai/salf

2026 VIENNA

UNIVERSITY OF TORONTO

## Motivation: Unified Multi-Sensor Simulation

- High-fidelity camera and LiDAR simulation is critical for autonomy testing and training
- NeRFs support flexible ray-based sensors, but are slow to train and render
- 3DGS is efficient, but requires additional approximations when moving beyond pinhole cameras

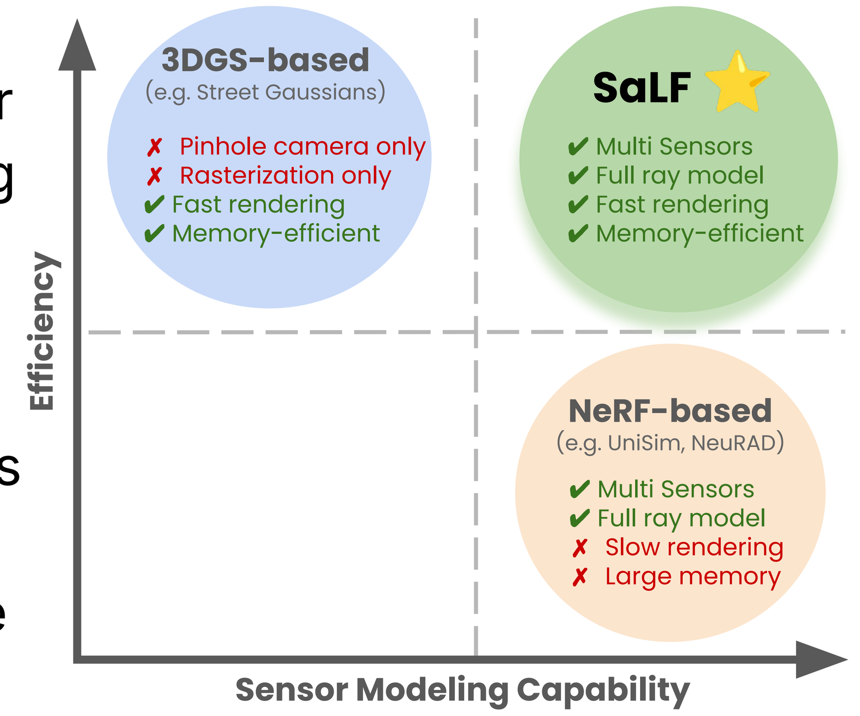


**Goal:** a unified representation supporting both **rasterization** and **ray-based rendering**

## Applications

- **Efficient and flexible multi-sensor simulation**

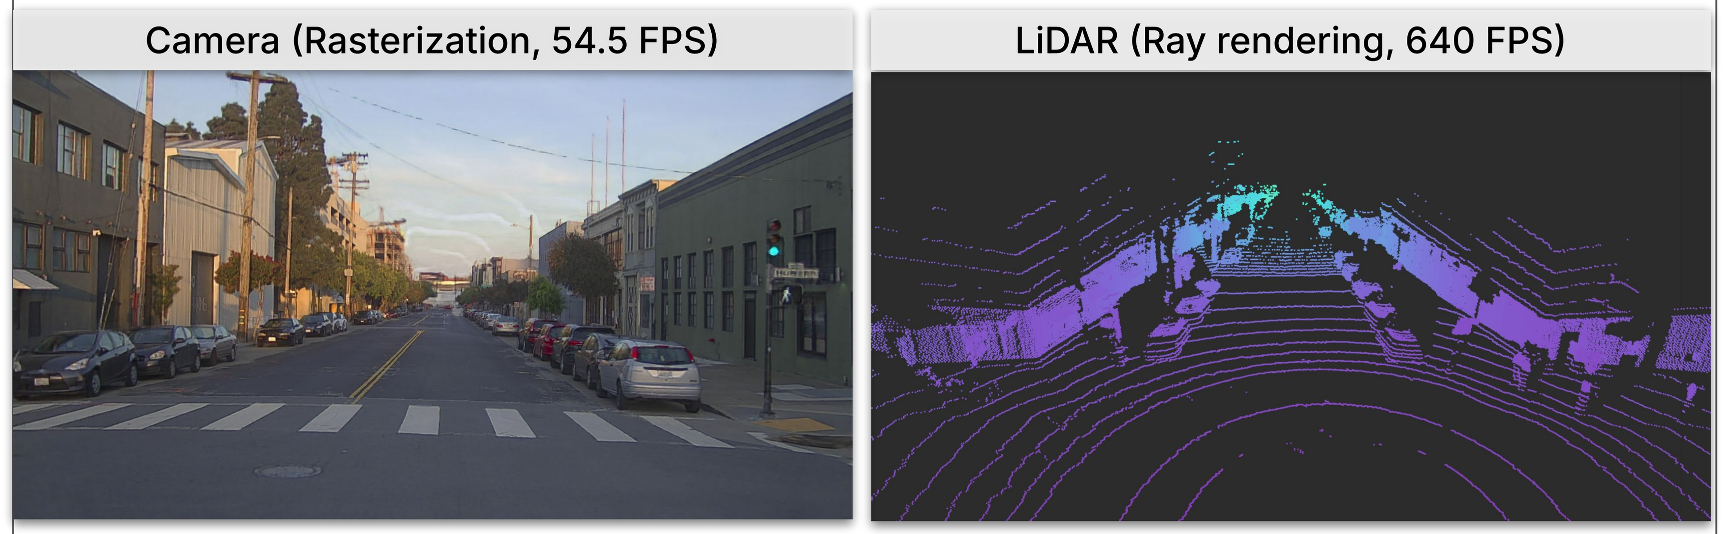


Multi-sensor rendering in real-time

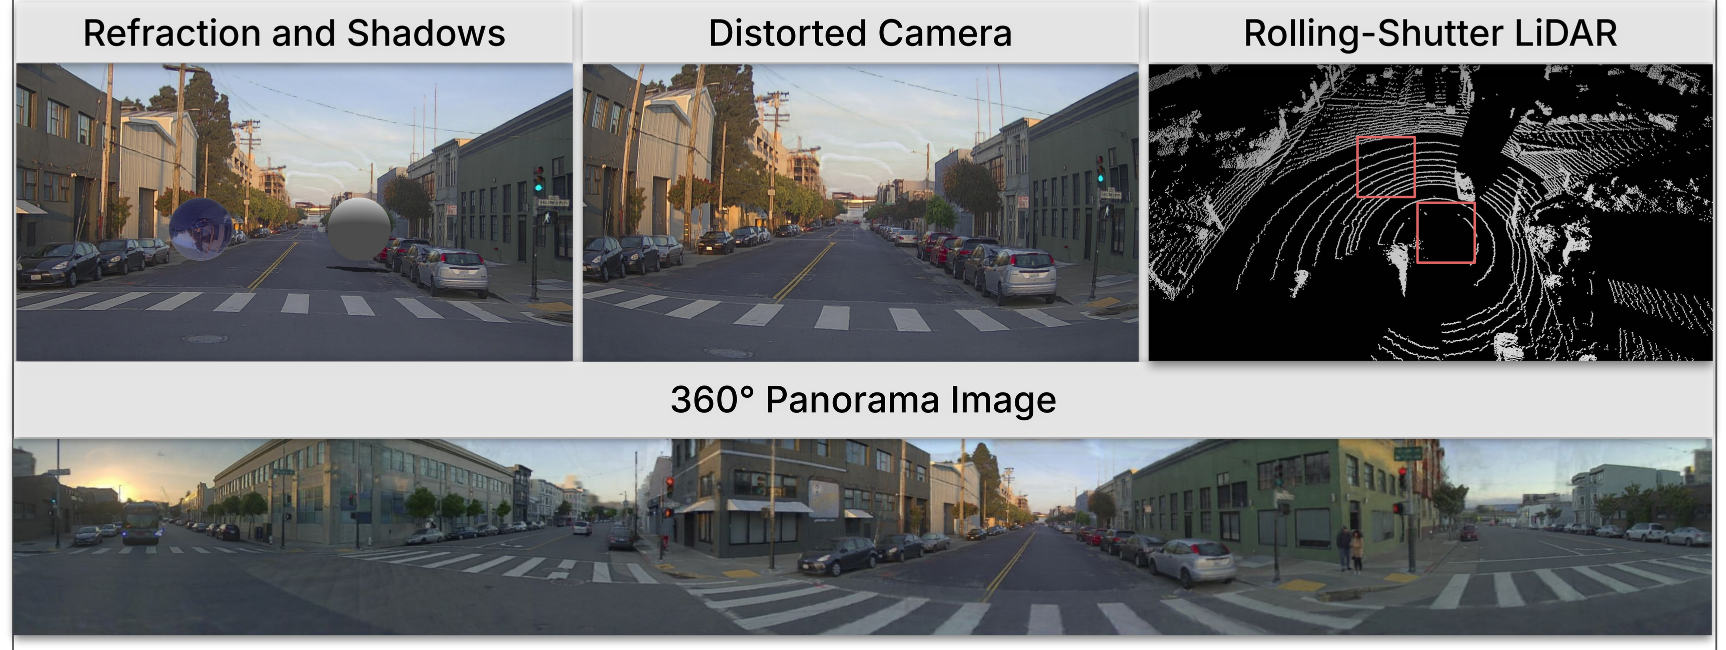


Advanced sensor modeling via ray rendering

## Representation

*SaLF* uses sparse local implicit voxel fields backed by an octree to enable efficient rasterization and ray rendering

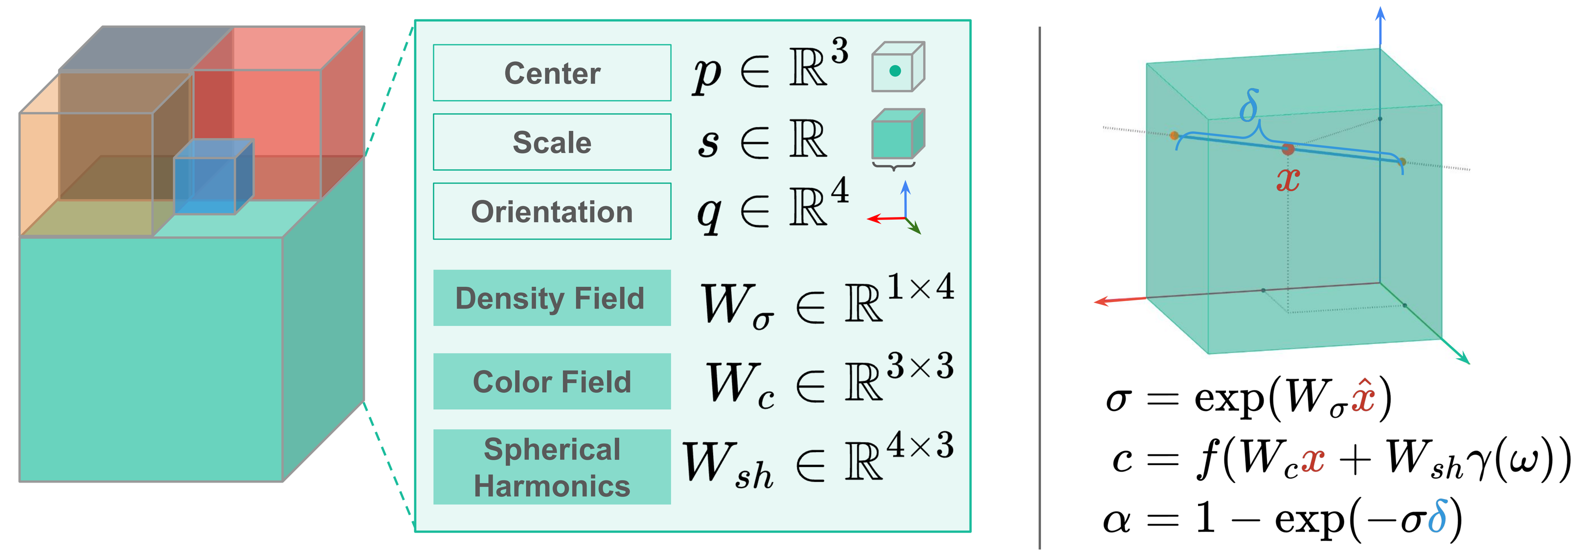


$$\sigma = \exp(W_\sigma \hat{x})$$
$$c = f(W_c x + W_{sh}\gamma(\omega))$$
$$\alpha = 1 - \exp(-\sigma\delta)$$

## Adaptive density control

Coarse-to-fine voxel densification subdivides high gradient regions and prunes empty space

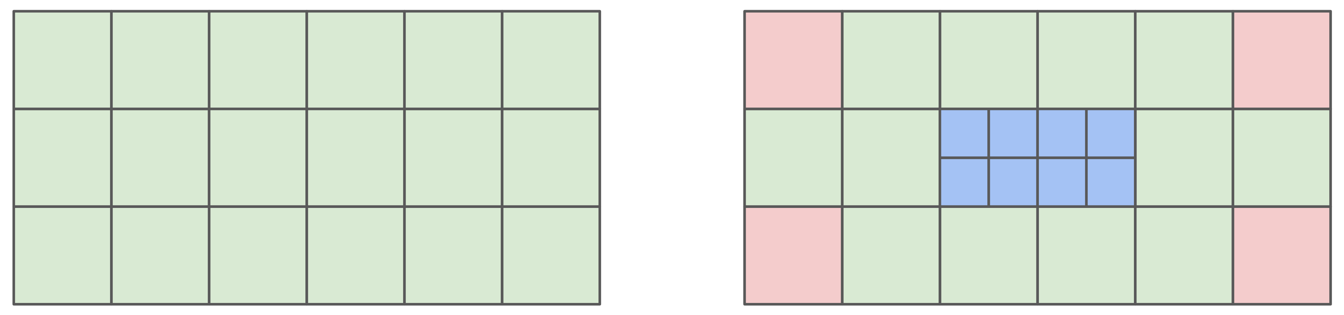

Dense Initialization | Prune and Split

## Rendering

One representation supports fast pinhole rasterization as well as flexible ray-based rendering

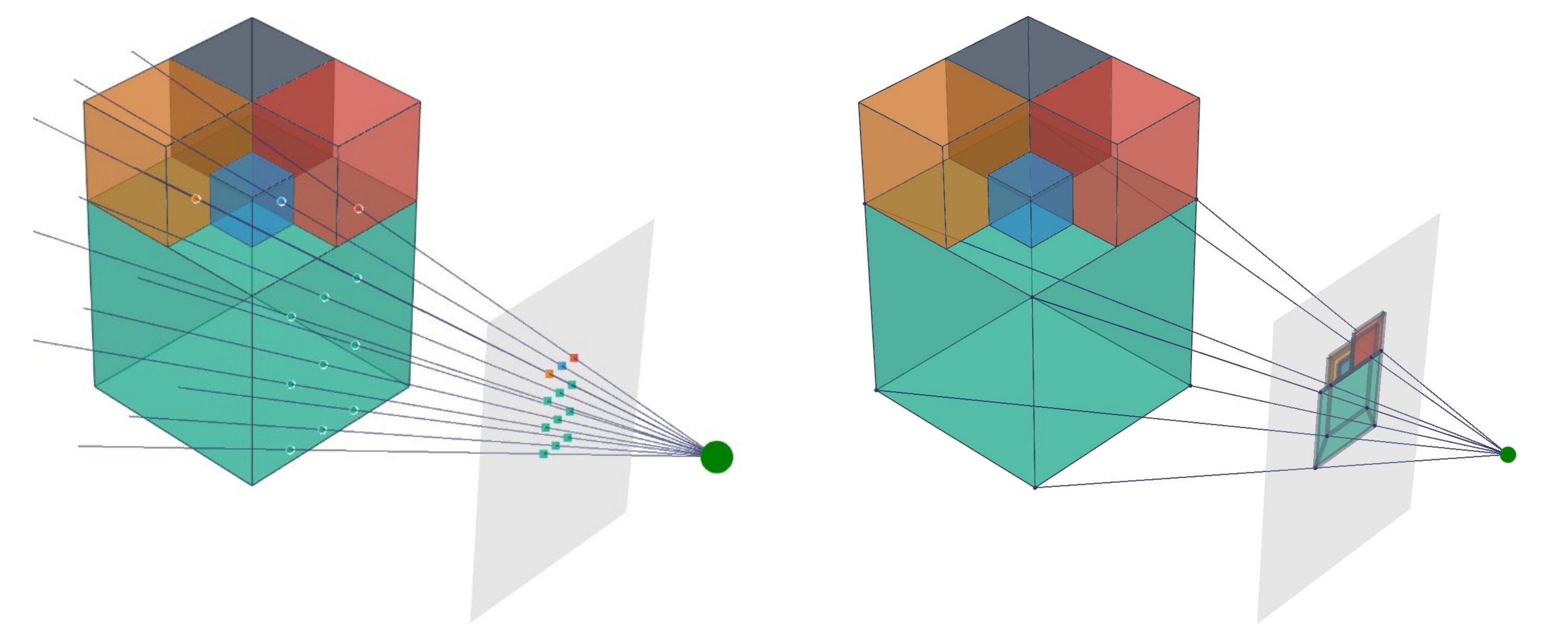

Octree-accelerated ray rendering | Tile-based rasterization

## Results

- **NVS comparison with SOTA**

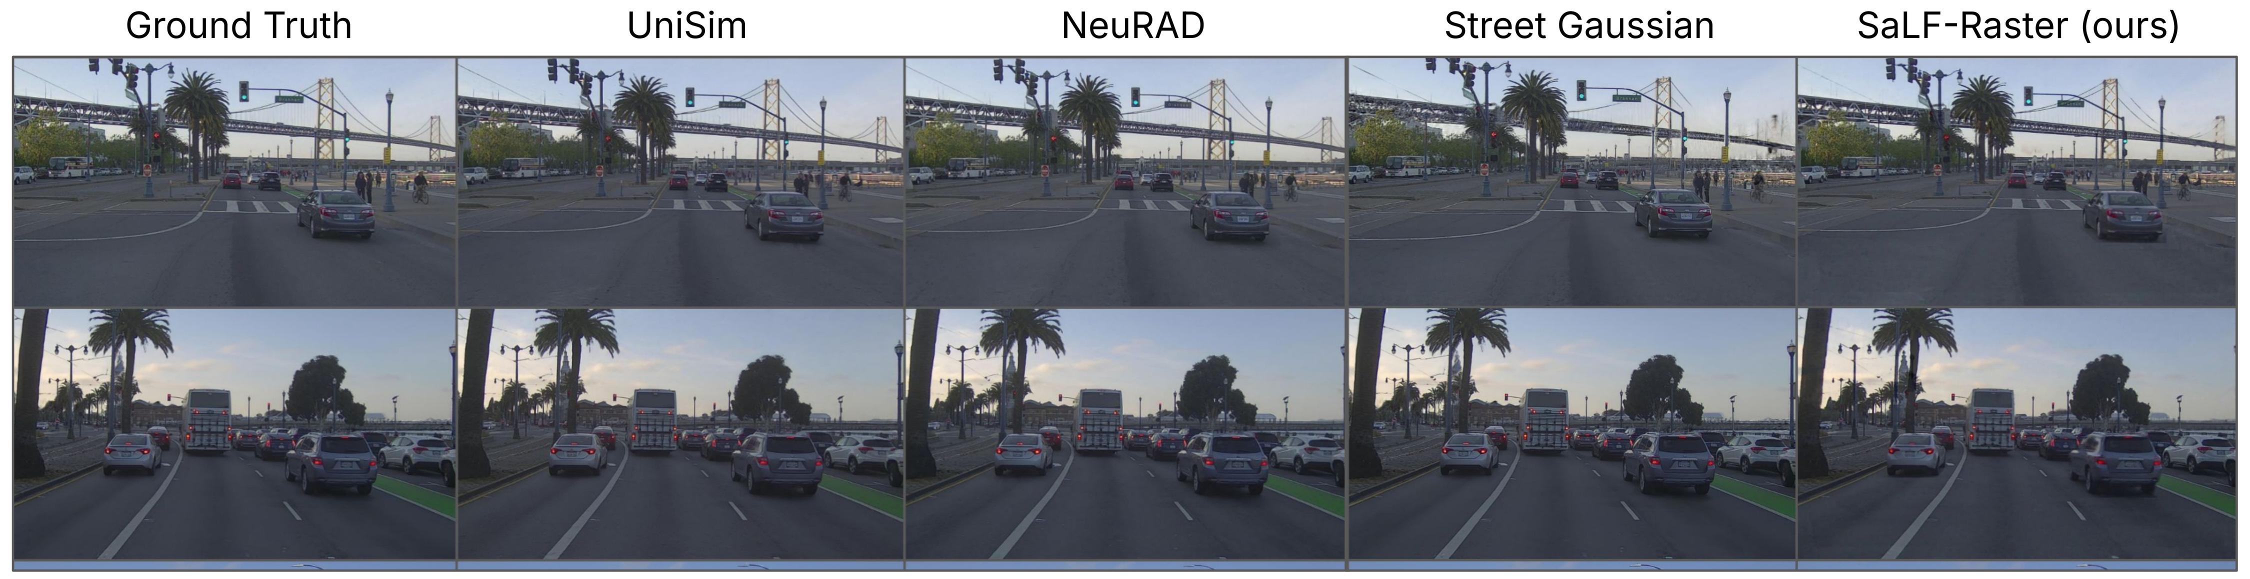


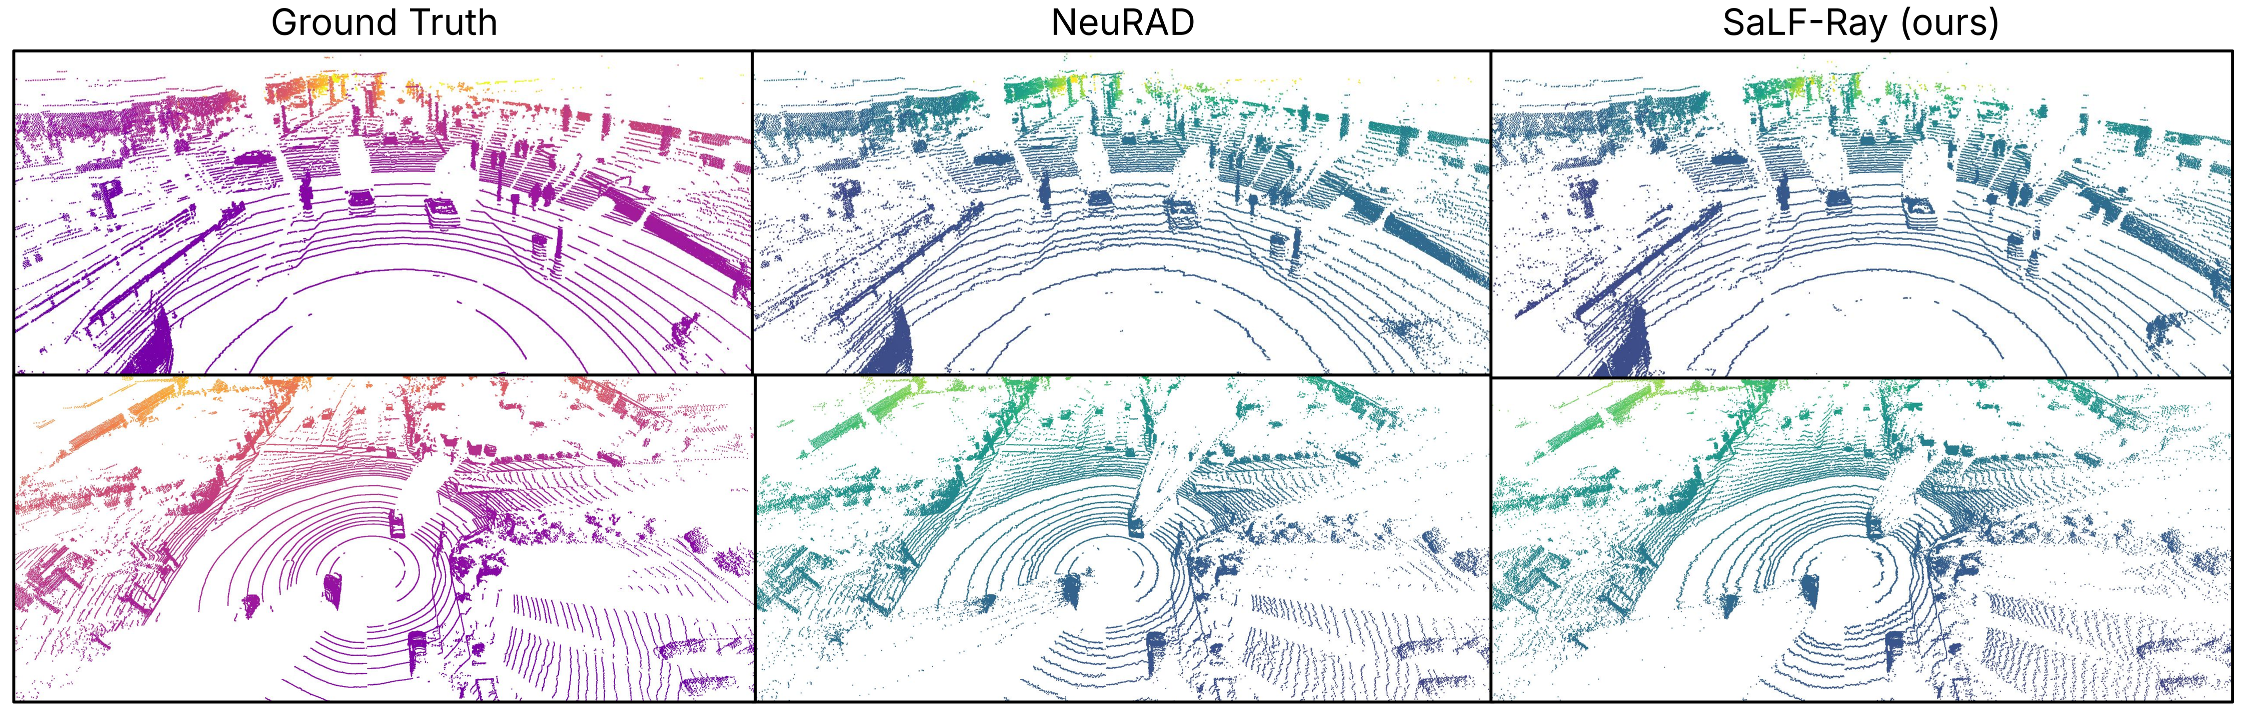


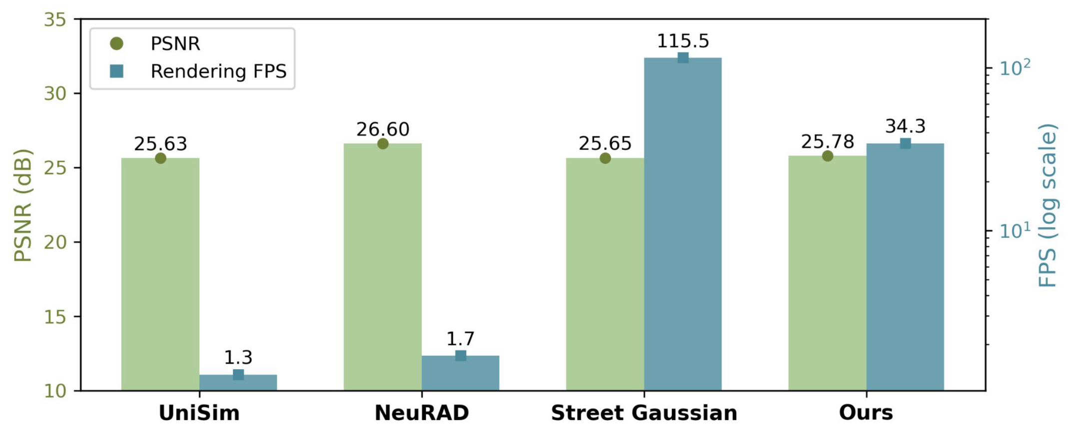


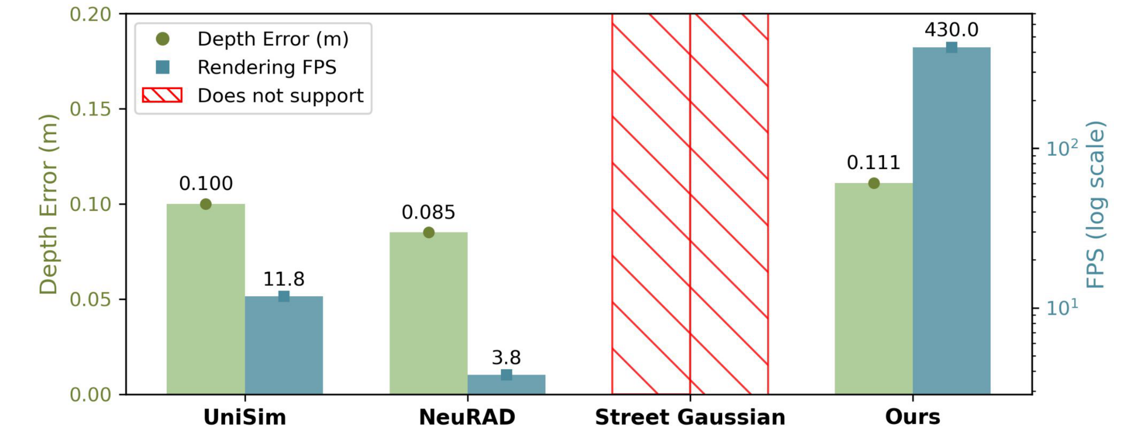


- **Ablation**

ABLATION STUDY ON SALF COMPONENTS.

| Models | PSNR↑ | SSIM↑ | LPIPS↓ |
|---|---|---|---|
| Ours | **25.48** | **0.744** | **0.373** |
| – Densification | 23.19 | 0.670 | 0.474 |
| – Field matrices | 25.11 | 0.735 | 0.386 |

- **Autonomy Evaluation**

DOWNSTREAM DOMAIN GAP EVALUATION.

| | Det. Agg. ↑ | Pred. ADE ↓ | Plan Cons. ↓ |
|---|---|---|---|
| UniSim | 0.74 | 0.63 | 0.99 |
| Ours | **0.78** | **0.52** | **0.83** |

- **Limitations:** (a) requires higher primitive count compared to 3DGS for similar quality, (b) doesn't support non-rigid and temporal changes in the representation